

# Evolution-Preserving Dense Trajectory Descriptors

Yang Wang, Vinh Tran, Minh Hoai

Stony Brook University, Stony Brook, NY 11794, USA

{wang33, tquangvinh, minhhoai}@cs.stonybrook.edu

## Abstract

*Recently Trajectory-pooled Deep-learning Descriptors were shown to achieve state-of-the-art human action recognition results on a number of datasets. This paper improves their performance by applying rank pooling to each trajectory, encoding the temporal evolution of deep learning features computed along the trajectory. This leads to Evolution-Preserving Trajectory (EPT) descriptors, a novel type of video descriptor that significantly outperforms Trajectory-pooled Deep-learning Descriptors. EPT descriptors are defined based on dense trajectories, and they provide complimentary benefits to video descriptors that are not based on trajectories. In particular, we show that the combination of EPT descriptors and VideoDarwin leads to state-of-the-art performance on Hollywood2 and UCF101 datasets.*

## 1. Introduction

The ability to recognize human actions in video has many potential applications in a wide range of fields. Recognizing human actions in realistic videos, however, is difficult due to various challenges such as background clutter, self occlusion, and viewpoint variation. An effective and efficient approach to handle these challenges is to use local visual space-time descriptors, an approach that does not require non-trivial pre-processing steps such as recognition and segmentation. In particular, one of the best types of video descriptor is Dense Trajectory Descriptors (DTD) [23, 24], which remain competitive even in the recent surge of deep-learning descriptors [18, 21, 25]. In fact, most recent human action recognition methods [1, 5, 9, 14, 25, 27, 28] propose to combine Dense Trajectory Descriptors with deep learning features to obtain the best results.

Human actions, by definition, are caused by the motion of humans, and Dense Trajectory Descriptors exploit this fact. The power and success of Dense Trajectory Descriptors can be attributed to their ability to extract local descriptors along humans' motion. The key idea is to densely sample feature points and track them in the video based

on optical flow. Typically, the feature points are tracked for 15 frames, leading to trajectories with a temporal span of 15 frames. DTD explicitly estimates the camera motion and prunes away trajectories from the background (i.e., Improved DTD [23]). It only retains trajectories from humans or objects of interest. For each trajectory, multiple descriptors are computed along the trajectory to capture shape, appearance, and motion information. In particular, each trajectory leads to four feature vectors: Histogram of Oriented Gradients [2], Histogram of Optical Flow (HOF), Motion Boundary Histogram (MBH), and Trajectory Geometry (TG). Recently, Wang et al. [25] propose to improve DTD by aggregating deep-learning descriptors (based on Two-stream CNN [18]) along each trajectory. This leads to the current state-of-the-art descriptors called Trajectory-pooled Deep-learning Descriptors (TDD).

However, TDD ignores the temporal evolution of appearance and motion along the trajectories, which turns out to be important for describing human actions. Recent work of Fernando et al. [4] has demonstrated the benefits of modeling video-wide temporal evolution, and they proposed a method called Rank Pooling for capturing temporal evolution. Unlike mean and max pooling methods that ignore the sequential order of video frames in a video, rank pooling provides a way to aggregate information while respecting the temporal order of the video frames. Fernando et al. [4] used rank pooling to encode the evolution of frame-based feature vectors, yielding impressive action recognition performance on a number of datasets.

In this paper, we propose Evolution-Preserving Trajectory (EPT) descriptors, ones that combine the benefits of dense trajectories, deep-learning features, and rank pooling. As depicted in Figure 1, EPT descriptors are based on dense trajectories, tagging along the apparent motion induced by human actions. EPT descriptors are based on deep-learning features, taking advantages of the feature representation produced by Convolutional Neural Networks [11, 13, 18] which have been trained on millions of images [17] and video frames [20]. EPT descriptors also capture the evolution of appearance and motion along the trajectories.

EPT descriptors require performing rank pooling of

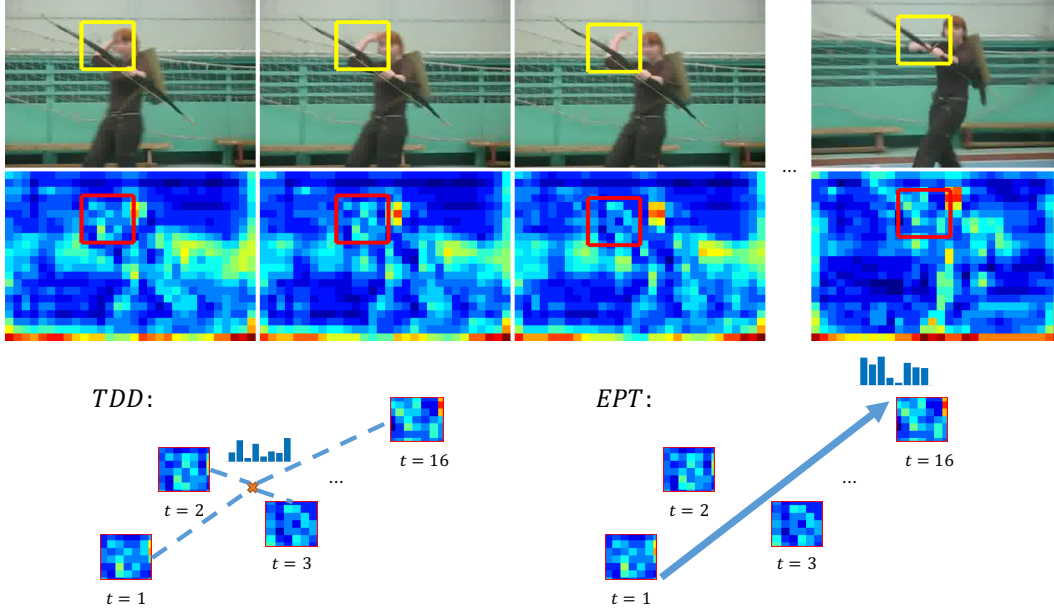


Figure 1. **Illustration of the Evolution-Preserving Trajectory (EPT) Descriptors.** Given a trajectory, we consider the volume encapsulating that trajectory, and extract a sequence of deep-learning features along the spatiotemporal volume. We propose EPT, a novel descriptor that preserves the temporal evolution of appearance and motion along the trajectories.

deep-learning features along densely populated trajectories. To avoid the prohibitive computational cost, we implement this procedure as follows. We first compute the video-wide convolution feature maps. We then use each trajectory’s coordinates to access its corresponding deep-learning features. To preserve the temporal evolution along the trajectories, we use approximate rank pooling instead of the exact rank pooling. Approximate rank pooling is linear operation, which is much more efficient than the optimizing the Ranking Support Vector Machine [4].

EPTs are local descriptors and they can be used in any standard recognition pipeline for action recognition. In this paper, we use EPT descriptors with Fisher Vector encoding [16] and Support Vector Machines [22]. We evaluate EPT descriptors on Hollywood2 and UCF101 datasets and find that they outperform the current state-of-the-art trajectories-based descriptors. When using EPT descriptors in conjunction with several pooling methods, the combined method outperforms the current state-of-the-art recognition performance on Hollywood 2 dataset.

## 2. Related Work

**Hand-crafted descriptors.** Feature descriptors are important for human action recognition, and much previous work has been devoted to the design of local features [3, 12, 23, 29], which are often robust to background clutter, self occlusion, and viewpoint variation in realistic video clips. There are various ways to extract informative

regions from a video clip. For instance, Space Time Interest Points [12] use 3D Harris corner detector, while Cuboid [3] uses temporal Gabor filters. In this paper, we choose the improved Dense Trajectories [23] due to its good performance. Each trajectory leads to four feature vectors: Trajectory Geometry, Histograms of Oriented Gradient [2], Histograms of Optical Flow, and Motion Boundary Histograms. These 4 descriptors are computed along the trajectory to capture information about shape, appearance, motion, and the change of motion, respectively.

**Deep-learning descriptors.** Deep learning has achieved great success in image based visual recognition tasks [6, 7, 11, 19]. There has been some successful attempts to develop deep architectures for action recognition [18, 21, 26] as well. Wang et al. [25] propose to improve dense trajectories descriptors by aggregating deep-learning descriptors (based on Two-stream CNN [18]) along each trajectory. This leads to the current state-of-the-art descriptors called Trajectory-pooled Deep-learning Descriptors (TDD).

**Rank Pooling.** In this paper, we propose to use rank pooling to preserve the temporal evolution along the trajectories. The concept of rank pooling is first introduced by Fernando et al. [4]. Their approach is referred to as Video-Darwin, which applies rank pooling to a sequence of video frames. Rank pooling has been applied to another extreme level of image pixels to produce dynamic images focusing on humans and objects of motion [1].

### 3. Evolution-Preserving Descriptors

In this section, we describe Evolution-Preserve Trajectory (EPT) descriptors, which are founded on dense trajectories, deep-learning features, and rank pooling.

#### 3.1. Improved Dense Trajectories

EPT descriptors are defined based on dense trajectories. To extract dense trajectories, we use the improved implementation of Wang and Schmid [23] instead of the original dense trajectories [24]. We customize the implementation of Wang and Schmid [23] with two modifications: (i) videos are normalized to have the height of 360 pixels, and (ii) frames are extracted at 25 fps. These modifications are added to standardize the feature extraction procedure across videos and datasets. They did not significantly alter the performance of the action recognition system. Note that each trajectory is specified as a sequence of  $L = 16$  points in the video space, corresponding to a temporal span of 15 frames. The coordinates of these points on the trajectory are later used to pool deep convolutional features.

#### 3.2. Evolution Preserving

The pipeline of computing deep-learning descriptors for trajectories is illustrated in Figure 2. Given a video, we first extract thousands of trajectories as local interest points in the video space, and densely apply a CNN to extract the convolutional feature maps. Subsequently, for each trajectory, we consider a spatiotemporal volume encapsulating that trajectory; the length of the volume is the same as the length of the trajectory, and the width and height of the volume are defined by the scale of the trajectory. Each trajectory-based volume is associated with a sequence of feature vectors  $f_1, \dots, f_L \in \mathbb{R}^d$ , with  $d$  being the number of CNN feature maps and  $L$  the length of the trajectory. The current state-of-the-art trajectory-based method, TDD [25], uses the average of  $f_1, \dots, f_L$  to represent the trajectory, i.e., to use  $u_0 = \frac{1}{L} \sum_{t=1}^L f_t$ . Average pooling is simple approach, but it ignores the temporal progression of the feature vectors along the trajectory. In order to capture temporal evolution of the feature vectors, we propose to replace average pooling with rank pooling. Rank pooling obtains a vector  $\mu^*$  to represent a sequence of feature vectors  $f_1, \dots, f_L$  by solving the following optimization problem:

$$\mu^* = \underset{\mu}{\operatorname{argmin}} \frac{\lambda}{2} \|\mu\|^2 + \sum_{i < j} \max\{0, 1 - \mu^T(F_j - F_i)\}$$

$$\text{where } F_t = \left\| \sum_{\tau=1}^t f_\tau \right\|_2.$$

This is the original formulation of rank pooling, which is based on Ranking Support Vector Machine. The param-

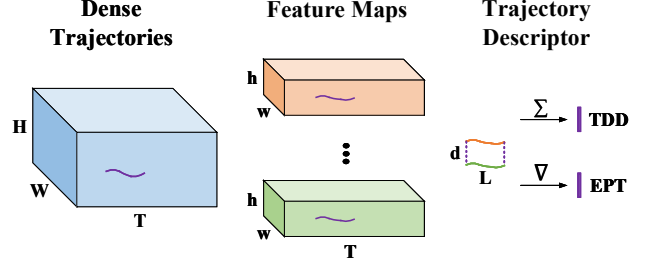


Figure 2. **Illustration of computing deep-learning descriptors for trajectories.** Given a video, We extract the dense trajectories and the convolutional feature maps. Each trajectory is associated with a sequence of feature vectors  $f_1, \dots, f_L \in \mathbb{R}^d$ , with  $d$  being the number of CNN feature maps and  $L$  the length of the trajectory. TDD applies Average Pooling to represent a trajectory. We propose to use Rank Pooling instead.

ter vector  $\mu$  is optimized to reflect the rank of the frames in the video. It therefore aggregates relevant information over time and encodes the temporal progression. This is a convex optimization problem. The first term of the objective is a commonly used quadratic regularizer. The second term is a Hinge-loss, penalizing the pairs of frames  $\{i, j\}$  that are incorrectly ranked. Note that a pair of frames  $\{i, j\}$  is correctly ranked only if the ranking scores of the  $L_2$ -normalized accumulative feature vectors  $\{F_i, F_j\}$  are separated by at least a unit margin, i.e.  $\mu^T(F_j - F_i) > 1$ .

However, the above formulation requires solving an optimization problem for each trajectory, and this can be computationally prohibitive because there can be hundred thousand of trajectories for each video clip. Therefore, we propose to use approximate rank pooling instead [1]. The approximate rank pooling is derived by initializing  $\mu$  as  $\vec{0}$  and updating  $\mu$  only once using gradient descent. As it turns out, the approximate rank pooling is essentially the average pooling of feature changes  $(F_j - F_i)$  over any time interval  $(1 \leq i < j \leq L)$ , which can be efficiently computed using a linear combination of the normalized accumulative feature vectors  $\{F_t\}_{t=1}^L$ :

$$u_1 = \sum_{1 \leq i < j \leq L} F_j - F_i = \sum_{t=1}^L (2t - L - 1) F_t. \quad (1)$$

#### 3.3. Convolutional Feature Maps

As shown in Figure 2, CNN feature maps are required to extract deep-learning descriptors. The process of CNN feature map computation is flexible with a variety of options, as listed in Table 1. It is up to us to decide which CNN model and which convolutional layer to be used to extract the feature maps. The size of the input video can also vary, because the convolutional layers do not require a specific size

Component	Options
model	spatial-stream VGG-M-2048
	temporal-stream VGG-M-2048
scale	$360p \times 480p$
	$240p \times 320p$
	$180p \times 240p$
layer	spatial: conv4, conv5
	temporal: conv3, conv4
normalization	in-voxel
	in-channel

Table 1. **Possible options for the convolutional feature maps.** EPT descriptors are flexible, they can be paired with both spatial- and temporal-stream CNNs. One can choose the size of the input images, the layer of the CNN models to be used for extracting feature maps. The output feature maps can be normalized with different normalization methods.

for the input as in the case of fully-connected layers. There are several different feature map normalization methods as well. In the following, we provide a detailed discussion of these options.

### 3.3.1 Feature Map Extraction

We first discuss and specify some of the options used in the process of extracting convolutional feature maps.

**Two-stream CNN.** In principle, any kind of CNN architecture can be used to extract convolutional feature maps. In our implementation, we follow the original TDD paper and use a Two-stream CNN provided by Wang et al. [25]. The model is trained on UCF-101 dataset and contains both a spatial and a temporal CNN. The spatial CNN is based on VGG-M-2048 model and fine-tuned with single RGB frames ( $224 \times 224 \times 3$ ); the temporal CNN has a similar structure, but its input is a volume of stacking optical flow fields ( $224 \times 224 \times 2F$ ,  $F = 10$  is the number of stacking flows).

For the purpose of extracting convolutional feature maps, we make two modifications to the original model. First, we make the model fully-convolutional by removing all the fully-connected layers. Second, we pad zero values to each convolution layer’s input with the size of the pad being half the size of the convolutional kernels. This is to avoid the boundary-shrinking effect of the convolutional kernels.

**Input Scale.** The size of the input RGB frames and optical flow images can vary, because we only use convolutional layers and convolutional layers do not require a fixed input size like the fully-connected layers. Having a larger input size means each neuron will look at a smaller neighborhood (perceptive field) and produce features at a finer scale.

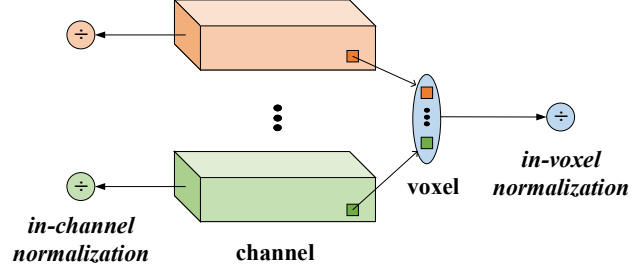


Figure 3. **Illustration of Feature Map Normalization.** Left: *in-channel* normalization is performed within each individual feature channel; Right: *in-voxel* normalization is performed within each individual voxel.

In our experiments, descriptors computed with larger input size usually has better performance. Combining descriptors at different input scales can improve the classification results.

**Spatial Feature Maps.** We extract frames from videos at 25fps and resize them into a predefined scale (e.g.,  $360p \times 480p$ ). Subsequently we feed individual frames into the pre-trained spatial CNN and take the output of the ‘conv4’ or ‘conv5’ layer as the spatial feature maps for each video, as suggested in the original TDD method [25].

**Temporal Feature Maps.** The process of computing the temporal feature maps is similar to the process of computing spatial feature maps, except we feed into the pre-trained temporal CNN with 10 consecutive optical flow images (with horizontal and vertical motion channels). The optical flow images are computed at 25fps using a GPU version of TVL1 algorithm, and subsequently resized into a predefined scale as well. For computing the temporal CNN descriptors, we used the ‘conv3’ or ‘conv4’ layer’s output as the temporal feature maps.

### 3.3.2 Feature Map Normalization

After the feature maps are extracted, we apply feature map normalization. Traditionally, normalization has been widely used for hand-crafted local descriptors such as HOG, HOF, and MBH. Here we discuss two methods that can be used to normalize convolutional feature maps: *in-channel* normalization and *in-voxel* normalization.

**In-channel normalization.** Figure 3 (left) illustrates the in-channel normalization method, which is defined as  $L_\infty$  normalization of each individual feature channel across the video’s spatiotemporal volume. This normalization is to ensure each feature channel ranges in the same interval, and thus contributes equally to final representation.

**In-voxel normalization.** Figure 3 (right) also illustrates the in-voxel normalization method. It is defined as  $L_\infty$  nor-

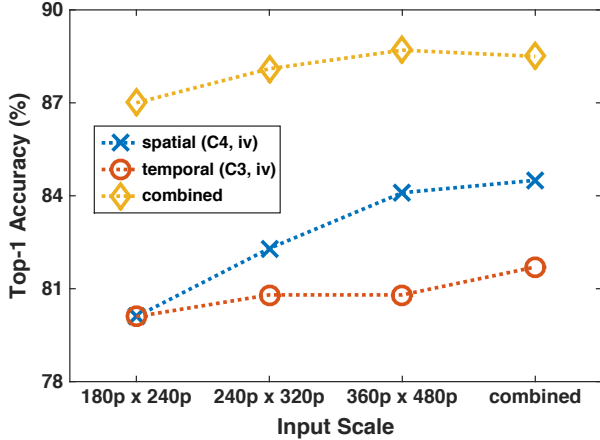


Figure 4. **The performance of EPT descriptors at different input scales.** On UCF101 dataset (split-1), the EPT descriptors computed at a larger input size usually outperforms those computed at a smaller input size. Combining descriptors at multiple scales achieve the best result, for both the spatial and temporal streams.

malization of each individual voxel across the feature channels. This ensures each voxel ranges in the same interval, and consequently makes equal contribution to the final representation.

### 3.4. Fisher Vector encoding

Finally, after computing the local descriptors, we use Fisher Vector [16] to represent an entire video. A Fisher Vector encodes both first and second order statistics between the feature descriptors and a Gaussian Mixture Model (GMM). In [23], Fisher Vector shows an improved performance over bag of features for action classification.

For each type of descriptors, we first sample a subset of 1,000,000 data points, and use PCA to de-correlate the descriptors and reduce the dimension to  $D$  ( $D = 64$  for CNN-based descriptors;  $D$  is half of the original dimension for iDT descriptors), and then train a GMM with  $K = 256$  mixtures. Finally, each video is represented with a  $2KD$ -dim Fisher Vector, which is subsequently power ( $\alpha = 0.5$ ) normalized and  $L_2$  normalized, as in [16, 23].

After encoding the Fisher Vectors, we train one-vs-all SVMs [22] for action classification. After learning the classifiers, given a test video, we compute the probability of each action by normalizing the scores across actions with a softmax function [28]. In order to fuse different descriptor types, we average their corresponding linear kernels, which is equivalent to concatenating their corresponding Fisher Vectors.

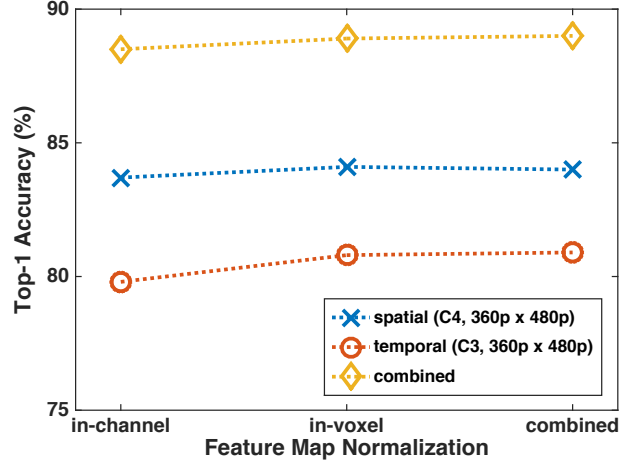


Figure 5. **Comparing different normalization methods.** This shows the performance of EPT descriptors on the UCF101 dataset (split-1). In-voxel normalization consistently outperforms in-channel normalization. Using both normalization methods does not lead to better performance than using in-voxel normalization alone. We suggest to use in-voxel normalization only.

## 4. Experiments

### 4.1. Datasets

We evaluate the performance of EPT descriptors using two datasets: Hollywood2 [15] and UCF101 [20]. The Hollywood2 dataset has 12 action classes and contains 1707 video clips collected from 69 different Hollywood movies. The videos are split into a training set of 823 videos and a testing set of 884 videos. The training and testing videos come from different movies. We augment the training set with horizontally flipped training videos, and report the mean Average Precision as the performance measure on Hollywood2. The UCF101 dataset contains 101 action categories with 13,320 realistic action videos collected from YouTube. Each category has at least 100 video clips. The dataset has three different training/test splits. We follow the evaluation scheme of the THUMOS13 challenge [10] and adopt the top-1 accuracy as the evaluation metric for each training/test split in UCF101.

### 4.2. Exploratory Studies

Table 1 lists 4 components that can be configured in our pipeline: *model*, *layer*, *scale*, *normalization*. However, the *model* and *layer* are not subject to change, because we use the same models and layers in the original TDD paper [25], to make the TDD baselines as strong as possible. As for the *scale* and *normalization*, we analyze their effect below.

**Input Scale.** We first investigate the performance of EPT descriptors computed at different input scales on the



UCF101 dataset (split-1). In this experiment, we fix the ‘conv4’ activations (after RELU) as the spatial feature maps and ‘conv3’ activations (after RELU) as the temporal feature maps. All feature maps are normalized using in-voxel normalization. We consider three input sizes,  $180p \times 240p$ ,  $240p \times 320p$ , and  $360p \times 480p$ , and the results are shown in Figure 4. As can be observed, the EPT descriptors computed at a larger input size performs better than those computed at a smaller input size. This might be attributed to having smaller receptive field. With a larger input size, each neuron looks at a smaller neighborhood (perceptive field) and consequently outputs the descriptors at a finer scale, leading to better performance. Of course, this is only true up to a certain extent. If the input size becomes extremely large, the neurons only observe tiny regions that are semantically meaningless. Also, a larger input size requires more computational power and memory capacity. In Figure 4, we also notice that combining descriptors at multiple scales can further improve the performance on both the spatial and temporal streams.

**Normalization method.** The normalization methods for convolutional feature maps is another important factor to investigate. In this experiment, the spatial feature maps (‘conv4’) and the temporal feature maps (‘conv3’) are both computed at input scale of  $360p \times 480p$ . The feature maps are either normalized using in-channel normalization or in-voxel normalization. Figure 5 compares the two normalization methods on the UCF101 dataset (split-1). In general, in-voxel normalization is slightly better than in-channel normalization for EPT descriptors. The reason might be that, with in-voxel normalization, each voxel is equally represented in the same interval and therefore making the rank pooling more effective in terms of discerning the changes among adjacent voxels. Figure 5 also suggests that the combination of the two normalization methods does not lead to significant improvement over the use of in-voxel normalization alone. Therefore, we only use in-voxel normalization in the following experiments.

**Chosen configuration.** Equipped with the knowledge learned from these exploratory experiments, we specify some options in Table 2 for the evaluation of EPT descriptors. In order to extract convolutional feature maps, we apply the Two-stream CNN provided by Wang et al. [25]. The ‘conv4’ activations (after RELU) of the spatial CNN are taken as the spatial feature maps. For temporal CNN, we consider the ‘conv3’ activations as the temporal feature maps. The input size of both RGB frames and optical flow images is set to be  $360p \times 480p$ . All convolutional feature maps are normalized using in-voxel normalization.

### 4.3. Evaluation of EPT Descriptors

We evaluate the performance of the proposed EPT descriptors on both Hollywood2 and UCF101 dataset (split-1),

Component	Used Configuration
model	Two-stream VGG-M-2048
scale	$360p \times 480p$
layer	spatial: conv4 temporal: conv3
normalization	in-voxel

Table 2. Chosen configuration for the CNN feature maps used by EPT descriptors. This configuration was determined based on our exploratory studies.

and the experiment results are summarized in Table 3. It’s worth noting that, as did in [4, 5], we apply rank-pooling in both forward and backward directions. The forward and backward EPT descriptors are then fused by concatenating their corresponding Fisher vectors. As can be observed, compared to the TDD descriptors, EPT descriptors significantly improve the recognition performance on both the spatial and temporal streams, which indicates the importance of preserving the temporal evolution of appearance and motion along the trajectories. The advantage is most evident on the spatial stream, where EPT outperforms TDD by 7% on both dataset (from 45.5% to 52.8% on Hollywood2; from 77.5% to 84.1% on UCF101). Figure 7 depicts the accuracy gap, the difference between the accuracy of spatial-TDD and the accuracy of spatial-EPT for 101 classes on the UCF101 dataset. As can be seen, most of the classes have a positive accuracy gap, indicating the effectiveness of rank pooling over average pooling. For some actions, the improvement can be as high as 40%. For the temporal stream, the EPT descriptors are also better than the TDD descriptors, even though the temporal CNN already takes as input a stack of optical flow images and models the temporal evolution of human motions.

In Figure 6, we show two video clips. Clip (a) is misclassified by the spatial TDD, but correctly classified by the proposed spatial EPT. Clip (b), on the other hand, is correctly classified by the spatial TDD, but misclassified by our spatial EPT. As can be seen, the amount of motion and tracked trajectories in Clip (b) is relatively small. Because the appearance change along the trajectory is small, the spatial EPT descriptor becomes close to zero, failing to describe the static scenes around the trajectory. This is a drawback of rank pooling over average pooling. Currently we address this issue by fusing EPT with TDD. In future work, we will seek a descriptor that encodes both the static appearance and the appearance change along the trajectories.

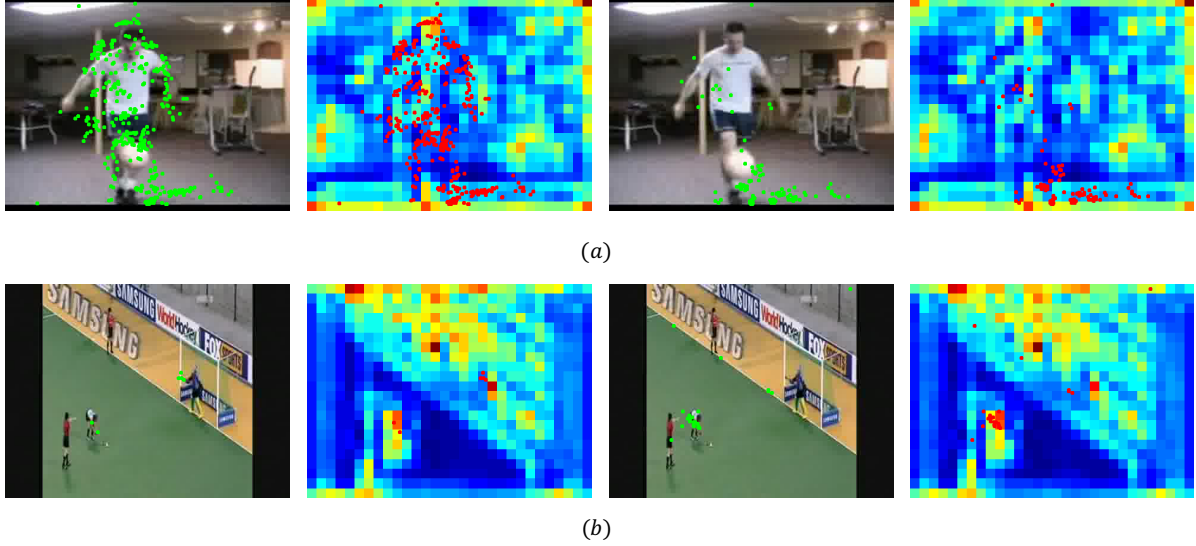


Figure 6. **Success case (a) and failure case (b) for the spatial EPT.** Clip (b) is correctly classified by the spatial TDD, but misclassified by the spatial EPT. The amount of motion and trajectories in Clip (b) is relatively small. Because the appearance change along the trajectory is small, the spatial EPT descriptor becomes close to zero, failing to describe the static scenes around the trajectory. This is a drawback of using rank pooling over average pooling. This problem will be addressed in our future work.

Dataset	Feature Maps	Trajectory-level Pooling		
		TDD	EPT	Improve
Hollywood2	spatial	45.5	<b>52.8</b>	7.3
	temporal	61.4	<b>63.3</b>	1.9
	2-stream	64.8	<b>67.1</b>	2.3
UCF101	spatial	77.5	<b>84.1</b>	6.6
	temporal	77.9	<b>80.8</b>	2.9
	2-stream	86.1	<b>88.7</b>	2.6

Table 3. **Comparing EPT and TDD descriptors on Hollywood2 and UCF101 (split-1).** EPT descriptors significantly improve the recognition performance on both the spatial and temporal streams, which indicates the importance of preserving the temporal evolution of appearance and motion along the trajectories. The advantage is most evident on the spatial stream.

#### 4.4. Using EPT descriptors with different video-level pooling methods

After verifying the advantage of EPT descriptors over TDD descriptors on the trajectory level, we further investigate the performance of EPT descriptors combined with different video-level pooling methods.

Each trajectory has a temporal span of 15 frames, and we assign each trajectory to its middle frame (the 8<sup>th</sup> frame). Each frame is therefore associated with a set of trajectories, and we can compute an unnormalized Fisher Vector for each frame. Consequently, a video can be represented as a sequence of frame-wise unnormalized Fisher Vectors

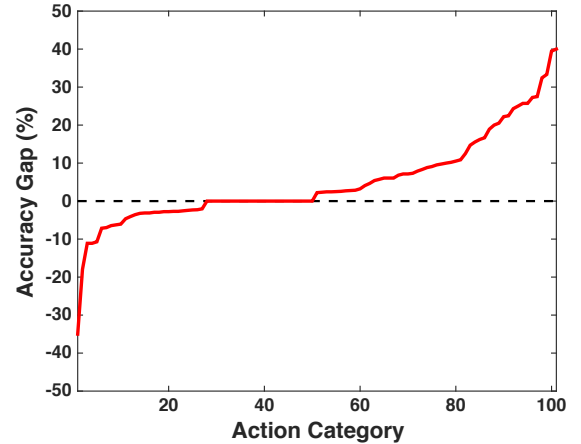


Figure 7. **Accuracy gap between spatial-EPT and spatial-TDD for 101 classes of the UCF101 dataset (split-1)** A positive gap indicates the improvement in classification accuracy from spatial-TDD to spatial-EPT. Most of the classes have positive accuracy gaps. For some actions, the improvement can be as high as 40%.

$\phi_1, \dots, \phi_T$ . In order to compress the information from the sequence into a single feature vector, we can apply many different video-level pooling methods. Up to now, we have been using average pooling to represent an entire video as  $\phi = \text{norm}(\sum_{t=1}^T \phi_t)$ . As aforementioned, we set the *norm* function to be power normalization ( $\alpha = 0.5$ ) followed by  $L_2$  normalization. We can also apply video-level rank pooling, i.e., VideoDarwin [4]). Aside from average pooling (AP) and rank pooling (RP), we also propose a sim-

Dataset	EPT	Video-level Pooling			
		AP	HAP	RP	HAP+RP
Hollywood2	spatial	52.8	54.1	55.9	<b>56.2</b>
	temporal	63.3	67.0	68.3	<b>68.8</b>
	2-stream	67.1	69.6	<b>71.4</b>	71.2
	+ DTD	72.5	74.6	<b>76.3</b>	<b>76.3</b>
UCF101	spatial	84.1	<b>84.5</b>	83.6	84.1
	temporal	80.8	82.6	82.1	<b>82.7</b>
	2-stream	88.7	<b>89.5</b>	88.8	89.2
	+ DTD	88.5	<b>89.5</b>	89.2	<b>89.5</b>

Table 4. **Comparing different video-level pooling methods for EPT descriptors.** Hierarchical average pooling (HAP) and rank pooling (RP) outperform average pooling (AP) video-wide, and their fusion can further improve the results. The improvement of video-wide rank pooling is more evident for the temporal stream. On Hollywood2, the fusion of EPT and DTD descriptors achieves significant improvement.

ple and effective pooling method called hierarchical average pooling (HAP). To obtain the final feature vector using hierarchical average pooling, we simply run a sliding window (size: 20, stride: 1) and perform average pooling and normalization within each window, and subsequently perform average pooling over the entire video again. We suggest to use HAP instead of AP, because HAP consistently outperforms AP in our experiments.

Table 4 compares the performance of EPT descriptors combined with different video-level pooling methods on Hollywood2 and UCF101 datasets. As can be observed, both the hierarchical average pooling (HAP) and the rank pooling (RP) outperform the ordinary average pooling (AP), and their fusion can further improve the classification results. Interestingly, the improvement of rank pooling over average pooling at the video level is more evident for the temporal stream, whereas at the trajectory level the improvement is higher for the spatial stream. On the Hollywood2 dataset, we notice that the fusion between EPT and DTD descriptors achieves significant improvement. On the UCF101 dataset, however, the combination with DTD does not lead to much better performance.

#### 4.5. Comparison to the state-of-the-art

In order to achieve the best recognition performance, we compute EPT descriptors with multiple convolutional layers (spatial: ‘conv4’+‘conv5’; temporal: ‘conv3’+‘conv4’) and input scales ( $180p \times 240p$ ,  $240p \times 320p$ ,  $360p \times 480p$ ). We also fuse EPT with DTD descriptors and apply the VideoDarwin technique. Table 5 compares our recognition performance with several other state-of-the-art methods on Hollywood2 and UCF101. On the Hollywood2 dataset, our

	Hollywood2	UCF101
DTD + FV [23]	64.7	85.9
SSD + RCS [8]	73.6	–
VideoDarwin [4]	73.7	–
Two Stream [18]	–	88.0
DIN + DTD [1]	–	89.1
C3D + DTD [21]	–	90.4
HRP + DTD [5]	76.7	91.4
TDD + DTD [25]	–	91.5
RNN-FV + DTD [14]	–	94.1
TSN [26]	–	<b>94.2</b>
Our method	<b>78.6</b>	92.3

Table 5. **Comparison with other state-of-the-art methods.** We compute EPT descriptors with multiple convolutional layers and input scales to preserve short-term evolution, and fuse them with DTD descriptors. We also apply VideoDarwin to capture long-term dynamics. On Hollywood2, our method outperforms previous state-of-the-art by 2%. On UCF101, we achieve performance on par with the Temporal Segment CNN, which is a very-deep version of the Two-stream CNN. In future work, we will use these new models to extract convolutional feature maps and seek further improvement.

method outperforms previous state-of-the-art by 2%. On the UCF101 dataset, we achieve performance on par with the Temporal Segment CNN [26], which is a very-deep version of the Two-stream CNN based on the VGG-VeryDeep-16 model [19] with three input modalities: RGB, optical flow, and warped optical flow. In future work, we will use these new models to extract convolutional feature maps and seek further improvement.

## 5. Conclusions

In this paper, we propose Evolution-Preserving Trajectory (EPT) descriptors, which integrate the benefits of dense trajectories, deep-learning features, and rank pooling. Deep architectures are utilized to extract discriminative convolutional feature maps, and rank pooling is applied to each trajectory, preserving the temporal evolution of appearance and motion along the trajectory. The EPT descriptor significantly outperforms the previous state-of-the-art trajectory-based descriptors. Combining EPT descriptors that preserve short-term evolution and the video-wide rank pooling that captures long-term dynamics, we are able to advance the state-of-the-art performance on Hollywood2 dataset.

## References

- [1] H. Bilen, B. Fernando, E. Gavves, A. Vedaldi, and S. Gould. Dynamic image networks for action recognition. In *Proc. CVPR*, 2016.



- [2] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *Proc. CVPR*, 2005.
- [3] P. Dollár, V. Rabaud, G. Cottrell, and S. Belongie. Behavior recognition via sparse spatio-temporal features. In *ICCV Workshop on Visual Surveillance & Performance Evaluation of Tracking and Surveillance*, 2005.
- [4] B. Fernando, E. Gavves, J. O. M., A. Ghodrati, and T. Tuytelaars. Modeling video evolution for action recognition. In *Proc. CVPR*, 2015.
- [5] B. Fernando, P. Anderson, M. Hutter, and S. Gould. Discriminative hierarchical rank pooling for activity recognition. In *Proc. CVPR*, 2016.
- [6] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proc. CVPR*, 2014.
- [7] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. *Proc. CVPR*, 2016.
- [8] M. Hoai and A. Zisserman. Thread-safe: Towards recognizing human actions across shot boundaries. In *Proc. ACCV*, 2014.
- [9] M. Jain, J. C. van Gemert, and C. G. Snoek. What do 15,000 object categories tell us about classifying and localizing actions? In *Proc. CVPR*, 2015.
- [10] Y. Jiang, J. Liu, A. R. Zamir, G. Toderici, I. Laptev, M. Shah, and R. Sukthankar. Thumos challenge: Action recognition with a large number of classes. In *ECCV Workshop*, 2014.
- [11] A. Krizhevsky, I. Sutskever, and G. Hinton. ImageNet classification with deep convolutional neural networks. In *NIPS*, 2012.
- [12] I. Laptev. On space-time interest points. *IJCV*, 64(2–3):107–123, 2005.
- [13] Y. LeCun, B. Boser, J. S. Denker, and D. Henderson. Back-propagation applied to handwritten zip code recognition. *Neural Computation*, 1(4):541–551, 1989.
- [14] G. Lev, G. Sadeh, B. Klein, and L. Wolf. Rnn fisher vectors for action recognition and image annotation. In *Proc. ECCV*, 2016.
- [15] M. Marszalek, I. Laptev, and C. Schmid. Actions in context. In *Proc. CVPR*, 2009.
- [16] F. Perronnin, J. Sánchez, and T. Mensink. Improving the fisher kernel for large-scale image classification. In *Proc. ECCV*, 2010.
- [17] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. Imagenet large scale visual recognition challenge. *IJCV*, 115(3):211–252, 2015.
- [18] K. Simonyan and A. Zisserman. Two-stream convolutional networks for action recognition in videos. In *NIPS*, 2014.
- [19] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *Proc. ICLR*, 2015.
- [20] K. Soomro, A. R. Zamir, and M. Shah. UCF101: A dataset of 101 human action classes from videos in the wild. Technical Report CRCV-TR-12-01, University of Central Florida, 2012.
- [21] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri. Learning spatiotemporal features with 3d convolutional networks. In *Proc. ICCV*. IEEE, 2015.
- [22] V. Vapnik. *Statistical Learning Theory*. Wiley, New York, NY, 1998.
- [23] H. Wang and C. Schmid. Action recognition with improved trajectories. In *Proc. ICCV*, 2013.
- [24] H. Wang, A. Klaser, C. Schmid, and C.-L. Liu. Action recognition by dense trajectories. In *Proc. CVPR*, 2011.
- [25] L. Wang, Y. Qiao, and X. Tang. Action recognition with trajectory-pooled deep-convolutional descriptors. In *Proc. CVPR*, 2015.
- [26] L. Wang, Y. Xiong, Z. Wang, Y. Qiao, D. Lin, X. Tang, and L. Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *Proc. ECCV*, pages 20–36. Springer, 2016.
- [27] P. Wang, Y. Cao, C. Shen, L. Liu, and H. T. Shen. Temporal pyramid pooling based convolutional neural networks for action recognition. *arXiv preprint arXiv:1503.01224*, 2015.
- [28] Y. Wang and M. Hoai. Improving human action recognition by non-action classification. In *Proc. CVPR*, 2016.
- [29] G. Willems, T. Tuytelaars, and L. V. Gool. An efficient dense and scale-invariant spatio-temporal interest point detector. In *Proc. ECCV*, 2008.